To: "Franck Chastagnol" <fchastagnol@youtube.com>
From: "Jim Schrempp" <j_schrempp@audiblemagic.com>
Cc: "v_ikezoye@audiblemagic.com" <v_ikezoye@audiblemagic.com>, "Jeremy Stern" <j_stern@audiblemagic.com>, "Doug Keislar" <doug@musclefish.com>, "Lou Kvitek" <l_kvitek@audiblemagic.com>
Bcc:
Received Date: 2006-10-19 21:39:00 CST
Subject: RE: Questions regarding reference fingerprints

Hi,

Yes, I understand that we should remove the human factor. We need to add some tests to the reference creation code that will warn you if the soundtrack does not have enough data to be useful. By the way, it doesn't take much to be useful, so most soundtracks will pass.

WRT your question, today our backend would probably let a very low complexity soundtrack to be added - we need to do some specific tests - perhaps you could send us a couple of samples. At the same time, it would only match another soundtrack that was the same. However, in your context different user videos might have the same low complexity soundtrack. We would identify them as the same. Of course, we are only testing a 20 second sample from - most likely - an offset of 7 seconds, plus or minus a bit. So the soundtracks would have to precisely match there. Still, in your realm some false positives are possible since the soundtracks might be identical. This won't be the case in the music and movie soundtrack arena.
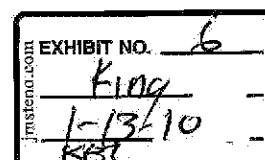
I see some ways to handle the situation.

a. After we identify a file in the YT database you could check other attributes of the user file. Length for instance.

b. We are working on a new fingerprint based on the image data itself. We're far enough along to know that it is workable. In November we should have something that we could demo outside of the company.

I hope this helps. Let me know your thoughts.

Best Regards,
Jim

Audible Magic Corporation
408.399.6405 x105

This message is intended for the use of the addressee only and may contain confidential information and trade secrets of Audible Magic. Unauthorized use or disclosure is prohibited.

From: Franck Chastagnol [mailto:fchastagnol@youtube.com]
Sent: Thursday, October 19, 2006 10:18 AM
To: Jim Schrempp
Cc: Doug Keislar; Vance Ikezoye (((((Vance Ikezoye)))); Frey Waid
Subject: Re: Questions regarding reference fingerprints

I agree that it would be good for us to build the functionality
into siggen to test the reference for silence. We haven't had this
problem with reference submitters before because they do not have
this kind of soundtrack to their videos. We will need to think
about what tests to add to siggen. I will get back to you on this.

please let us know. it would be excellent to have this included in
the tool

In the mean time, since your take down content will be reviewed by
a human before submission, can the human flag videos with silent
soundtracks and divert those away from the YT reference database?

some content will be reviewed but not all of it.
and anyway, i would not want to rely on human, even one well
trained... they make mistakes :)

regarding my question in the previous email, what do you think ?

3. AM works very, very well with non-music. We do require some
"interest" in the audio. For instance, if the soundtrack was just
a continuous 60 hertz hum, then it would match any other
soundtrack with just a continuous 60 hertz hum.

To be clear, will your back-end detect such a reference
fingerprint is not appropriate ?

thanks,
franck

Regards,
Jim

Audible Magic Corporation
408.399.6405 x105

———

From: Franck Chastagnol [mailto:fchastagnol@youtube.com]
Sent: Wednesday, October 18, 2006 7:29 PM
To: Jim Schrempp
Cc: 'Frey Waid'; 'Lou Kvitek'; 'Doug Keislar'; 'Nancy Kang'
Subject: Re: Questions regarding reference fingerprints

Hi Jim,

Thanks for the answers.
I have a few follow up questions regarding 2. and 3.
See below

Thanks,
Franck

> 1. For your application a reference fingerprint should come from
an offset of 0 and a duration of 60. You may remember that we had
discussed doing a more expensive search of the YT database, but
that was decided to be out of scope.

2. Our id requires 20 seconds of contiguous audio with less than
one second of silence. Our backend will scan the 60 second
reference fingerprint you send us to find an appropriate 20
seconds to use. This brings up a good point - we need to provide
you with some feedback on reference fingerprints that do not meet
this criteria.

As opposed to receiving a feed from you with a list of invalid ref
fingerprints we submitted,

for us it would be ideal if the amSigGen tool could detect a
fingerprint is not valid at the time of its generation.
Could the functionality be added to amSigGen ?

3. AM works very, very well with non-music. We do require some
"interest" in the audio. For instance, if the soundtrack was just
a continuous 60 hertz hum, then it would match any other
soundtrack with just a continuous 60 hertz hum.

To be clear, will your back-end detect such a reference
fingerprint is not appropriate ?

4. I think the siggen application does support .wav files for
input. If you get some error, please let me know.

Jim

Audible Magic Corporation
408.399.6405 x105

This message is intended for the use of the addressee only
and may contain confidential information and trade secrets
of Audible Magic. Unauthorized use or disclosure is
prohibited.

—————

From: Franck Chastagnol [mailto:fchastagnol@youtube.com]
Sent: Wednesday, October 18, 2006 1:58 PM
To: Jim Schrempp
Cc: Frey Waid
Subject: Re: Questions regarding reference fingerprints

Added question 4. below
Franck

On Oct 17, 2006, at 4:44 PM, Franck Chastagnol wrote:

Jim,

Who would be the engineer we should talk to regarding reference
fingerprints ?

See below our list of questions:
1. Optimum length for reference fingerprint ? minimum length ?
maximum length ?
1a. Approximate size per minute of reference fingerprint?
1b. Recommended to reference against entire audio track (<= 10
minutes)?
2. How to prevent a "vanilla" reference fingerprint to be uploaded
causing thousands of match ? For instance a blank reference
fingerprints causing all videos having little or no audio track to
be found as match.
3. How well does AM work with non-music?

4. Could the amSigGenLinux tool take as input a .wav file as
opposed to a mp3 ?

Thanks,
Franck

_____